

What Is Claimed Is:

1. A computer-implemented method of adding a new node to a network multicast group, with a specified group membership status, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:
 - selecting a minimum spanning tree of the network;
 - selecting the new node as the current node;
 - examining the routing tree membership statuses of nodes that are linked to the current node by links included in the minimum spanning tree;
 - until said examining is halted, selecting a peer node of the current node as the current node and repeating said examining;
 - halting the examining when a final node is examined if:
 - the final node is a Full member of the routing tree; or
 - the final node is a SendOnly member of the routing tree and the specified group membership status of the new node is SendOnly; and
 - for each node from the new node to the final node, setting the routing tree membership status of the node equal to the specified group membership status of the new node.
- 20
2. The method of claim 1, further comprising:
 - maintaining a queue for storing network nodes for selection as current node;
 - wherein a first peer of a current node is added to said queue if:
 - 25
said first peer is coupled to the current node by a link included in the minimum spanning tree;
 - said first peer is not a Full member of the routing tree; and

the routing tree membership status of said first peer and the specified group membership status of the new node are not both SendOnly.

- 5 3. The method of claim 2, further comprising:
determining if said queue is empty if:
the specified group membership status of the new node is Full; and
the routing tree membership status of said first peer is SendOnly.
- 10 4. The method of claim 2, further comprising:
halting the examining if said queue is empty.
- 15 5. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of adding a new node to a network multicast group, with a specified group membership status, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:
selecting a minimum spanning tree of the network;
selecting the new node as the current node;
20 examining the routing tree membership statuses of nodes that are linked to the current node by links included in the minimum spanning tree;
until said examining is halted, selecting a peer node of the current node as the current node and repeating said examining;
halting the examining when a final node is examined if:
25 the final node is a Full member of the routing tree; or
 the final node is a SendOnly member of the routing tree and the specified group membership status of the new node is SendOnly; and

for each node from the new node to the final node, setting the routing tree membership status of the node equal to the specified group membership status of the new node.

5 6. A computer-implemented method of adding a first node to a multicast group of network nodes, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

- (a) receiving a first request to include a first network node in a multicast group as one of a Full member and a SendOnly member;
- 10 (b) setting a GroupStatus of the first node according to the request, wherein said GroupStatus indicates a membership status in the multicast group;
- (c) selecting a minimum spanning tree of the network;
- (d) selecting the first node as the current node;
- 15 (e) selecting a peer node of the current node, wherein a TreeStatus of the selected peer has not been examined since the first request was received, wherein said TreeStatus indicates a membership status in the routing tree;
- (f) performing one or more of the following examinations:
 - (f1) determining if said TreeStatus of the selected peer is Full;
 - 20 (f2) determining if said TreeStatus of the selected peer is SendOnly and said GroupStatus of the current node is SendOnly; and
 - (f3) determining if a network link coupling the current node to the selected peer is part of the selected minimum spanning tree;
- (g) repeating (e) – (g) until one of:
 - (g1) at least one peer of the current node has been examined; and
 - (g2) one of said (f1) and said (f2) determinations succeed;

(h) if neither of said (f1) and said (f2) determinations has succeeded, setting a peer of the current node as the current node;

(i) repeating (e) – (h) until one of:

(i1) all nodes in the routing tree have been examined; and

5 (i2) one of said (f1) and said (f2) determinations succeed;

(j) for each node in the minimum spanning tree, from the new node to the last peer examined, setting a TreeStatus of the node equal to said GroupStatus of the new node.

10 7. The method of claim 6, further comprising, after (d): comparing said GroupStatus of the first node to said TreeStatus of the first node.

15 8. The method of claim 6, further comprising: maintaining a queue in which to queue nodes for selection as the current node.

9. The method of claim 8, further comprising, if said (f3) determination succeeds:

20 adding the selected peer to said queue.

10. The method of claim 9, further comprising, if said (f3) determination succeeds:

25 if said TreeStatus of the selected peer is SendOnly and said GroupStatus of the new node is Full, determining if said queue is empty.

11. A computer readable medium storing instructions that, when

executed by a computer, cause the computer to perform a method of adding a first node to a multicast group of network nodes, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

5 (a) receiving a first request to include a first network node in a
multicast group as one of a Full member and a SendOnly member;

10 (b) setting a GroupStatus of the first node according to the request,
wherein said GroupStatus indicates a membership status in the multicast group;

15 (c) selecting a minimum spanning tree of the network;

(d) selecting the first node as the current node;

(e) selecting a peer node of the current node, wherein a TreeStatus of
the selected peer has not been examined since the first request was received,
wherein said TreeStatus indicates a membership status in the routing tree;

20 (f) performing one or more of the following examinations:

(f1) determining if said TreeStatus of the selected peer is Full;

(f2) determining if said TreeStatus of the selected peer is
SendOnly and said GroupStatus of the current node is SendOnly; and

(f3) determining if a network link coupling the current node to
the selected peer is part of the selected minimum spanning tree;

25 (g) repeating (e) – (g) until one of:

(g1) at least one peer of the current node has been examined;
and

(g2) one of said (f1) and said (f2) determinations succeed;

(h) if at least one peer of the current node has been examined, setting a
peer of the current node as the current node;

(i) repeating (e) – (h) until one of:

(i1) all nodes in the routing tree have been examined; and

(i2) one of said (f1) and said (f2) determinations succeed;

(j) for each node in the minimum spanning tree, from the new node to the last peer examined, setting a TreeStatus of the node equal to said GroupStatus of the new node.

5

12. A computer-implemented method of adding a new node to a network multicast group, with a specified group membership status, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

10 identifying a minimum spanning tree of the network;

selecting the new node as the current node;

until a final node having a routing tree membership status greater than or equal to the specified group membership status of the new node is identified, repeating:

15 examining the routing tree membership statuses of peer nodes of the current node; and

selecting as current node a peer node of the current node that is coupled to the current node by a link included in the minimum spanning tree; and

20 setting the routing tree membership status of each node in the minimum spanning tree, from the new node to the final node, to the specified group membership status of the new node.

13. The method of claim 12, wherein a node's routing tree membership status and group membership status are each one of the following, from lesser status to greater status: non-member, SendOnly, Full.

14. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of adding a new node to a network multicast group, with a specified group membership status, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

5 identifying a minimum spanning tree of the network;

 selecting the new node as the current node;

 until a final node having a routing tree membership status greater than or equal to the specified group membership status of the new node is identified,

10 repeating:

 examining the routing tree membership statuses of peer nodes of the current node; and

 selecting as current node a peer node of the current node that is coupled to the current node by a link included in the minimum spanning tree; and

15 setting the routing tree membership status of each node in the minimum spanning tree, from the new node to the final node, to the specified group membership status of the new node.

20 15. A computer-implemented method of removing a first node from a network multicast group, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

 queueing the first node in a queue;

25 until the queue is empty, repeating the following, in order:

 (a) removing the most recently queued node to serve as the current node;

(b) returning to (a) if the group membership status of the current node is Full;

(c) identifying a number of local ports of the current node that are on;

5 (d) returning to (a) if the number is greater than one;

(e) if the number of local ports that are on is equal to zero:

(e1) for each peer node having a local port to the current node on, turning off said peer node's local port to the current node and adding said peer node to the queue; and

10 (e2) setting the routing tree membership status of the current node to None; and

(f) if the number of local ports that are on is equal to one:

(f1) on a sole peer node coupled to the one local port, turning off the sole peer node's local port to the current node if the sole peer's local port to the current node is on;

15 (f2) adding the sole peer node to the queue;

(f3) if zero peer nodes have local ports to the current node on and the group membership status of the current node is None:

20 turning off the one local port of the current node that is on; and

setting the routing tree membership status of the current node to None; and

25 (f4) otherwise, setting the routing tree membership status of the current node to SendOnly.

16. The method of claim 15, further comprising, prior to said

repeating:

setting the group membership status of the first node to one of None and SendOnly.

5 17. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of removing a first node from a network multicast group, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

10 queueing the first node in a queue;

until the queue is empty, repeating the following, in order:

 (a) removing the most recently queued node to serve as the current node;

 (b) returning to (a) if the group membership status of the

15 current node is Full;

 (c) identifying a number of local ports of the current node that are on;

 (d) returning to (a) if the number is greater than one;

 (e) if the number of local ports that are on is equal to zero:

20 (e1) for each peer node having a local port to the current node on, turning off said peer node's local port to the current node and adding said peer node to the queue; and

 (e2) setting the routing tree membership status of the current node to None; and

25 (f) if the number of local ports that are on is equal to one:

 (f1) on the sole peer node coupled to the one local port, turning off the sole peer node's local port to the current node if the

sole peer's local port to the current node is on;

(f2) adding the sole peer node to the queue;

(f3) if zero peer nodes have local ports to the current node on and the group membership status of the current node is

5 None:

turning off the one local port of the current node that is on; and

setting the routing tree membership status of the current node to None; and

10 (f4) otherwise, setting the routing tree membership status of the current node to SendOnly.

18. A computer-implemented method of removing a first node from a multicast group of network nodes, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

15 receiving a first request to remove a first network node from membership in a multicast group, wherein the first node was one of a Full member and a SendOnly member of the multicast group;

20 setting a GroupStatus of the first node to one of None and SendOnly, wherein said GroupStatus indicates a membership status in the multicast group;

queueing the first node in a queue;

until the queue is empty, repeating:

(a) dequeuing a node from the queue to be the current node;

25 (b) determining if the GroupStatus of the current node is Full;

(c) determining a number of local ports of the current node that are on;

(d) if the number is equal to zero:

(d1) for each peer of the current node with a local port to the current node turned on:

(d1') setting the local port of the peer to off; and

(d1") adding the peer to the queue; and

5 (d2) setting a TreeStatus of the current node to None, wherein said TreeStatus indicates a membership status in the routing tree; and

(e) if the number is equal to one:

10 (e1) on the one peer coupled to the one local port of the current node, setting the local port of the one peer to the current node to off;

(e2) adding the one peer to the queue;

15 (e3) if the GroupStatus of the current node is None and zero peers of the current node have a local port to the current node on:

(e3') turning off the one local port of the current node; and

(e3") setting the TreeStatus of the current node to None;

and

20 (e4) otherwise, setting the TreeStatus of the current node to SendOnly.

19. The method of claim 18, wherein said (a) comprises:

20 dequeuing the node most recently added to the queue.

20. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of removing a first node from a multicast group of network nodes, wherein members of a corresponding routing tree are configured to route multicast messages among members of the group, the method comprising:

receiving a first request to remove a first network node from membership

in a multicast group, wherein the first node was one of a Full member and a SendOnly member of the multicast group;

setting a GroupStatus of the first node to one of None and SendOnly, wherein said GroupStatus indicates a membership status in the multicast group;

5 queuing the first node in a queue;

until the queue is empty, repeating:

(a) dequeuing a node from the queue to be the current node;

(b) determining if the GroupStatus of the current node is Full;

(c) determining a number of local ports of the current node that are on;

10 (d) if the number is equal to zero:

(d1) for each peer of the current node with a local port to the current node turned on:

(d1') setting the local port of the peer to off; and

(d1'') adding the peer to the queue; and

15 (d2) setting a TreeStatus of the current node to None, wherein said TreeStatus indicates a membership status in the routing tree; and

(e) if the number is equal to one:

(e1) on the one peer coupled to the one local port of the current node, setting the local port of the one peer to the current node to off;

20 (e2) adding the one peer to the queue;

(e3) if the GroupStatus of the current node is None and zero peers of the current node have a local port to the current node on:

(e3') turning off the one local port of the current node; and

(e3'') setting the TreeStatus of the current node to None;

25 and

(e4) otherwise, setting the TreeStatus of the current node to SendOnly.

21. A system for managing membership in a multicast group and a corresponding routing tree for routing multicast messages within the multicast group, the apparatus comprising:

5 a network node coupling the apparatus to a network;

 a subnet administrator configured to receive requests to change the membership of the multicast group;

 a subnet manager configured to update network nodes' routing tables when the routing tree is modified in response to a change in membership of the

10 multicast group; and

 a subnet management coordinator configured to:

 make a non-member into a Full or SendOnly member of the multicast group;

 make a Full or SendOnly member into a non-member of the

15 multicast group; and

 update the membership of the routing tree in response to a change in the membership of the multicast group.

22. The system of claim 21, wherein said subnet management

20 coordinator makes a non-member into a Full or SendOnly member of the multicast group by:

 setting the group membership status of the non-member to the group membership status specified in a request that was received to make the non-member a member of the multicast group;

25 identifying a minimum spanning tree of the network;

 selecting the non-member as the current node;

 until a final node having a routing tree membership status greater than or

equal to the group membership status of the non-member is identified, repeating:

examining the routing tree membership statuses of peer nodes of

the current node; and

selecting as current node a peer node of the current node that is

5 coupled to the current node by a link included in the minimum spanning tree;

setting the routing tree membership status of each node in the minimum spanning tree, from the non-member to the final node, to the specified group membership status of the new node.

10

23. The system of claim 22, wherein a node's routing tree membership status and group membership status are each one of the following, from lesser status to greater status: non-member, SendOnly, Full.

15

24. The system of claim 21, wherein said subnet management coordinator makes a Full or SendOnly member into a non-member of the multicast group by:

queueing the member in a queue;

until the queue is empty, repeating the following, in order:

20

(a) removing the most recently queued member to serve as the current node;

(b) returning to (a) if the group membership status of the current node is Full;

25

(c) identifying a number of local ports of the current node that are on;

(d) returning to (a) if the number is greater than one;

(e) if the number is equal to zero:

(e1) for each peer of the current node that has a local port to the current node on, turning off said peer's local port to the current node and adding said peer to the queue; and

(e2) setting the routing tree membership status of the

current node to non-member; and

(f) if the number is equal to one:

(f1) on the one peer coupled to the one local port,

turning off the peer's local port to the current node;

(f2) adding the one peer to the queue;

(f3) if zero peers have local ports to the current node on and the group membership status of the current node is non-member:

turning off the one local port of the current node that is on; and

setting the routing tree membership status of the current node to non-member; and

(f4) otherwise, setting the routing tree membership of the current node to SendOnly.

20 25. The system of claim 21, wherein said network node is one of a
channel adapter and a network switch.